arXiv:1410.4240v1 [cs.DS] 15 Oct 2014

Online Network Design Algorithms via Hierarchical Decompositions*

Seeun Umboh†

Abstract

We develop a new approach for online network design and obtain improved competitive ratios for several problems. Our approach gives natural deterministic algorithms and simple analyses. At the heart of our work is a novel application of embeddings into hierarchically well-separated trees (HSTs) to the analysis of online network design algorithms — we charge the cost of the algorithm to the cost of the optimal solution on any HST embedding of the terminals. This analysis technique is widely applicable to many problems and gives a unified framework for online network design.

In a sense, our work brings together two of the main approaches to online network design. The first uses greedy-like algorithms and analyzes them using dual-fitting. The second uses tree embeddings — embed the entire graph into a tree at the beginning and then solve the problem on the tree — and results in randomized $O(\log n)$-competitive algorithms, where n is the total number of vertices in the graph. Our approach uses deterministic greedy-like algorithms but analyzes them via HST embeddings of the terminals. Our proofs are simpler as we do not need to carefully construct dual solutions and we get $O(\log k)$ competitive ratios, where k is the number of terminals.

In this paper, we apply our approach to obtain deterministic $O(\log k)$-competitive online algorithms for the following problems.

1. Steiner network with edge duplication. Previously, only a randomized $O(\log n)$-competitive algorithm was known.
2. Rent-or-buy. Previously, only deterministic $O(\log^2 k)$-competitive and randomized $O(\log k)$-competitive algorithms by Awerbuch, Azar and Bartal (*Theoretical Computer Science* 2004) were known.
3. Connected facility location. Previously, only a randomized $O(\log^2 k)$-competitive algorithm of San Felice, Williamson and Lee (*LATIN* 2014) was known.
4. Prize-collecting Steiner forest. We match the competitive ratio first achieved by Qian and Williamson (*ICALP* 2011) and give a simpler analysis.

Our competitive ratios are optimal up to constant factors as these problems capture the online Steiner tree problem which has a lower bound of $\Omega(\log k)$.

1 Introduction

We study network design problems in the online model. In a network design problem, we are given a graph with edge costs and connectivity requirements. The goal is to find a minimum-cost subgraph satisfying the requirements. In the online model, requests in the form of terminals and connectivity requirements arrive one-by-one, and the goal is to maintain a minimum-cost subgraph satisfying all arrived requirements. The challenge in the online model is that decisions are irrevocable: once an edge is added to the subgraph, it cannot be removed later on. For instance, the classic Steiner tree problem asks for a minimum-cost subgraph connecting a given set of terminals. In the online Steiner tree problem [14], terminals arrive one-by-one and the algorithm maintains a subgraph connecting all terminals so far. In the following, we use n to denote the number of vertices in the underlying graph, and k the number of terminals.

*This work was done while the author was visiting Microsoft Research and the University of Washington. Supported in part by NSF Award CCF-1320854.

†University of Wisconsin – Madison. seeun@cs.wisc.edu.

Our main contribution is a novel application of embeddings into hierarchically well-separated trees (HSTs) [4, 5, 8] as a tool for analyzing online network design algorithms. In this paper, we apply this technique to several problems and obtain natural algorithms that improve upon previous work. Our approach also gives a simple unified analysis across these different problems.

At a high level, our work brings together two disparate lines of work in online network design. The first designs greedy-like algorithms and analyzes them via dual-fitting. The idea is to grow a dual solution in tandem with the algorithm such that at any step, the dual accounts for the cost of the algorithm, mirroring the primal-dual schema in the offline setting. Then one shows that the dual is feasible after appropriate scaling. Berman and Coulston [6] used this approach to give a $O(\log k)$-competitive algorithm for online Steiner forest which was later extended by Qian and Williamson [17] to the more general online constrained forest and prize-collecting Steiner forest problems.

The other line of work is based on tree embeddings. Awerbuch and Azar [2] observed that for many online network design problems, one can first probabilistically embed the entire input graph into a tree [4, 5, 8] before the requests arrive and then (essentially) solve the online problem on the tree. In the analysis, one shows that for any choice of tree embedding, the cost of the resulting solution can be charged to the cost of the optimal solution on the tree. For problems that are easily solved on trees, this results in a randomized algorithm whose competitive ratio is, in expectation over the probabilistic embedding, at most $O(\log n)$.

We combine these two approaches in the following way: we use greedy-like algorithms but analyze them using HST embeddings. In particular, we develop a charging scheme showing that for *any* expanding embedding of the terminals into a HST T, the cost of the algorithm is at most a constant times the cost of the optimal solution on the tree $\mathrm{OPT}(T)$. We emphasize that the embedding only exists in the analysis, not in the algorithm. This shows that the competitive ratio of the algorithm is bounded by

$$O(1) \cdot \min_{T \in \mathcal{T}} \frac{\mathrm{OPT}(T)}{\mathrm{OPT}}, \tag{1}$$

where $\mathcal{T}$ is the set of expanding HST embeddings of the terminals and OPT is the cost of the optimal solution on the input graph. The bound (1) is at most the expected distortion of the embedding of the terminals into HSTs which is $O(\log k)$ [4, 5, 8]. However, since (1) optimizes with respect to a particular subgraph, namely OPT, the bound can be much better depending on the instance. Unlike previous work using tree embeddings, we get better competitive ratios (deterministic $O(\log k)$ instead of randomized $O(\log n)$). Furthermore, our algorithms are natural greedy-like algorithms that do not need to know the entire underlying graph upfront.

Our work highlights an interesting connection between HST embeddings and dual-fitting. We can interpret our charging scheme in terms of dual-fitting — given any expanding HST embedding of the terminals, we build a dual solution that is feasible for the HST and charge against it. These dual solutions can be highly infeasible for the original graph, but averaging over the probabilistic embeddings of [8] gives a dual solution that is feasible after scaling by the embedding distortion. However, our approach differs from dual-fitting in the usual sense — we charge against *multiple* dual solutions simultaneously (one per embedding), and the dual solutions are not built with respect to the original metric but with respect to the HST embeddings of the terminals. In a sense, our work uses HST embeddings as a black-box to generate good dual solutions. Compared to the usual dual-fitting approach, we hide the complexity of the dual construction within the HST embedding, allowing us to give simpler algorithm descriptions since the algorithm no longer needs to take the dual into account. Furthermore, we get a more streamlined analysis and a unified approach to different problems. In some cases, we also get a tighter upper bound (e.g. see footnote 4 in Section 4.2).

1.1 Our Results

In this paper, we illustrate our technique on a variety of online network design problems. In all of these problems, there is an underlying graph $G = (V, E)$ with edge lengths $d(u, v)$. We assume w.l.o.g. that G is a complete graph, the edge lengths satisfy the triangle inequality, i.e. (V, d) is a metric space, and that the minimum edge length is 1. Initially, the algorithm does not know G; at each time step online, it only knows the submetric over the arrived terminals.

For a (multi-)graph H, we define its *cost* $c(H)$ to be the sum of lengths of edges in H. We use n to denote the number of vertices $|V|$ in the graph G, and k the number of terminals that arrive online. We remark that our competitive ratios below are optimal up to constants as these problems capture the online Steiner tree problem which has a $\Omega(\log k)$ lower bound [14].

Steiner problems. In the *online Steiner tree* problem, the algorithm is given a root terminal r at the beginning. Terminals i arrive online one-by-one and the algorithm maintains a subgraph H connecting terminals to the root. In the *online Steiner forest problem*, terminal pairs (s_i, t_i) arrive online one-by-one and the algorithm maintains a subgraph H in which each terminal pair is connected. In the *online Steiner network problem with edge duplication*, each (s_i, t_i) pair comes with a requirement R_i, and the algorithm maintains a multigraph H which contains R_i edge-disjoint (s_i, t_i)-paths. Note that allowing H to be a multigraph means that the algorithm is allowed to buy multiple copies of an edge. For brevity, we will simply call this the *online Steiner network problem*. The goal in these problems is to maintain a minimum-cost (multi-)graph H satisfying the respective requirements subject to the constraint that once an edge is added to H, it cannot be removed later on.

Previously, deterministic $O(\log k)$-competitive algorithms were known for the online Steiner tree and Steiner forest problems [14, 6]. However, the best algorithm (as far as we know) for the online Steiner network problem is to use tree embeddings and yields a randomized $O(\log n)$ competitive ratio. Our first result closes the gap between the known competitive ratios of these problems.

Theorem 1.1. *There is a deterministic $O(\log k)$-competitive algorithm for the online Steiner network problem (with edge duplication).*

We remark that our approach also gives a simpler analysis of the Berman-Coulston online Steiner forest algorithm [6].

Rent-or-buy. The rent-or-buy problem generalizes the Steiner forest problem. The algorithm is allowed to either *rent* or *buy* edges in order to satisfy a request. Buying an edge costs M times more than renting, but once an edge is bought, it can be used for free in the future. On the other hand, a rented edge can only be used once; future terminals have to either re-rent it or buy it in order to use it.

More formally, in the *online rent-or-buy problem*, the algorithm is also given a parameter $M \geq 0$. The algorithm maintains a subgraph H of bought edges; when a terminal pair (s_i, t_i) arrives, the algorithm buys zero or more edges and rents edges Q_i such that $H \cup Q_i$ connects s_i and t_i. Both rent and buy decisions are irrevocable — edges cannot be removed from H later on and Q_i is fixed after the i-th step. The total cost of the algorithm is $Mc(H) + \sum_i c(Q_i)$. In the *online single-source rent-or-buy problem*, the algorithm is also given a root terminal r in advance; terminals i arrive online and i has to be connected to r in the subgraph $H \cup Q_i$.

Previously, Awerbuch, Azar and Bartal [3][1] gave a randomized $O(\log k)$-competitive algorithm as well as a deterministic $O(\log^2 k)$-competitive algorithm and posed the existence of a deterministic $O(\log k)$-competitive algorithm as an open problem. Our second result resolves this positively.

Theorem 1.2. *There is a deterministic $O(\log k)$-competitive algorithm for the online rent-or-buy problem.*

Connected facility location. In the *online connected facility location problem*, we call terminals *clients*. At the beginning, the algorithm is given a parameter $M \geq 0$, a set of *facilities* $F \subseteq V$ and facility opening costs f_x for each facility $x \in F$. There is also a designated root facility $r \in F$ with zero opening cost. The algorithm maintains a set of *open* facilities $F' \subseteq F$ and a Steiner tree H connecting F' and r. At any online step, the algorithm knows the submetric over the arrived clients and the facilities. When a client i arrives, the algorithm may open a new facility, and then assigns i to some open facility $\sigma(i) \in F'$. All decisions are irrevocable — edges cannot be removed from H later on, clients cannot be reassigned and opened facilities cannot be closed. The total cost of the algorithm is $\sum_{z \in F'} f_z + Mc(H) + \sum_i d(i, \sigma(i))$. When $M = 0$ (i.e. the

[1] They called it the *network connectivity leasing problem*.

facilities need not be connected), this is called the *online facility location problem* [16]. For these problems, we use k to denote the number of clients.

This problem was recently proposed by San Felice, Williamson and Lee [18] and they gave a randomized $O(\log^2 k)$-competitive algorithm[2]. Our third result improves on their work.

Theorem 1.3. *There is a deterministic $O(\log k)$-competitive algorithm for the online connected facility location problem.*

Prize-collecting Steiner forest. In the *online prize-collecting Steiner forest* problem, each terminal pair (s_i, t_i) arrives with a penalty π_i and the algorithm either pays the penalty or connects the pair. The total cost of the algorithm is $c(H) + \sum_{i \in P} \pi_i$ where H is the subgraph maintained by the algorithm and $i \in P$ if the algorithm paid the penalty for (s_i, t_i). Note that penalties are irrevocable: once the algorithm decides to pay the penalty π_i, the penalty does not get removed from the algorithm's cost even if later on H ends up connecting s_i and t_i.

For this problem, we obtain the same deterministic $O(\log k)$ competitive ratio first achieved by Qian and Williamson [17] but our charging scheme is simpler than their dual-fitting analysis.

Theorem 1.4 ([17])**.** *There is a deterministic $O(\log k)$-competitive algorithm for the online prize-collecting Steiner forest problem.*

Our Techniques. For each problem, we design a natural greedy algorithm and show that the cost of the algorithm is at most a constant times the cost of the optimal solution on any HST embedding of the terminals. The key structural property of a HST embedding T is that each edge of T defines a cut $C \subseteq X$ whose diameter is proportional to the length of the edge (see Definition 2), so we can express $\mathrm{OPT}(T)$ in terms of contributions from bounded-diameter cuts. In the analysis, we charge the cost of the algorithm to these cuts by dividing the cost of the algorithm among terminals and charging the cost share of each terminal to a bounded-diameter cut of T containing it. We then argue that since the algorithm is greedy and the cuts have bounded diameter, each cut receives a charge of at most a constant times its contribution to $\mathrm{OPT}(T)$. Thus, the cost of the algorithm is at most $O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of the terminals. Since [8] gives a probabilistic HST embedding of the terminals with expected distortion $O(\log k)$, there exists a HST embedding T^* such that $\mathrm{OPT}(T^*) \leq O(\log k)\,\mathrm{OPT}$. This gives us the desired $O(\log k)$ bound on the competitive ratio of the algorithm. Note that we do not need any particular property of the embedding of [8]. It is only used to prove the existence of T^*.

1.2 Related Work

There is a long history of research on network design in the offline setting. Agrawal, Klein and Ravi [1] gave a primal-dual 2-approximation algorithm for the Steiner forest problem. This was later extended by Goemans and Williamson [11] to give a 2-approximation for the more general constrained forest and the prize-collecting Steiner tree problems. In a breakthrough result, Jain [15] gave a 2-approximation algorithm using iterative LP rounding for the Steiner network problem. We remark that Jain's result holds even when edge duplication is not allowed. Swamy and Kumar [19] gave a primal-dual 4.55-approximation algorithm for the single-source rent-or-buy problem and a 8.55-approximation algorithm for the connected facility location problem. Later, Gupta et al. [13] improved the constants using an elegant sample-and augment approach.

Online network design was first studied by Imase and Waxman [14] who showed that the greedy algorithm is $O(\log k)$-competitive for the online Steiner tree problem and gave a matching lower bound. Awerbuch, Azar and Bartal [3] showed that a greedy algorithm is $O(\log^2 k)$-competitive for the online Steiner forest problem; later, Berman and Coulston [6] gave a $O(\log k)$-competitive algorithm. Recently, Qian and Williamson [17] extended the approach of [6] to give an $O(\log k)$-competitive algorithm for the online constrained forest and prize-collecting Steiner forest problems. As far as we know, there is no previous work on the online Steiner

[2] At the time of submission, the author heard from San Felice (personal communication) that he independently obtained a randomized $O(\log k)$-competitive algorithm recently.

network problem with edge duplication. When edge duplication is disallowed, unlike the offline setting, the problem becomes much harder. Gupta et al. [12] showed a lower bound of $\Omega(\min\{k, \log n\})$ and an upper bound of $\tilde{O}(R_{\max} \log^3 n)$, where $R_{\max}$ is the maximum requirement.

The online rent-or-buy problem was first considered by Awerbuch, Azar and Bartal [3], and they gave a deterministic $O(\log^2 k)$-competitive algorithm and a randomized $O(\log k)$-competitive algorithm. The online connected facility location problem was recently proposed by San Felice et al [18] and they presented a randomized $O(\log^2 k)$-competitive algorithm based on the offline connected facility location algorithm of Eisenbrand et al. [7]. Surprisingly, Fotakis [10] showed that there is a lower bound of $\Omega(\frac{\log k}{\log \log k})$ for the special case of the facility location problem ($M = 0$) even when the underlying metric space is a HST.

For tree embeddings, Bartal showed that any graph can be probabilistically embedded into HSTs with $O(\log^2 n)$ expected distortion [4] and subsequently improved this to $O(\log n \log \log n)$ [5]. Fakcharoenphol et al. [8] achieved the tight $O(\log n)$ bound.

Roadmap. We start with the necessary HST embedding definitions and results in Section 2. Next, we warm up by illustrating our techniques to analyze the greedy algorithm for the Steiner tree problem in Section 3. In Section 4, we apply our analysis framework to the Berman-Coulston algorithm for online Steiner forest algorithm [6] (yielding a simpler analysis) and show that this essentially allows us to reduce the online Steiner network instance to several online Steiner forest instances.

We use the same high-level approach for the rent-or-buy, connected facility location and prize-collecting Steiner problems as they share very similar cost structures. The key ideas are illustrated using the single-source rent-or-buy problem in Section 5.1. Then we extend these ideas to the multi-source setting in Section 5.2 and the connected facility location problem in Section 6. The objective of the prize-collecting Steiner tree problem is closely related to that of the single-source rent-or-buy problem. In Section 7, we show how a straightforward adaptation of the single-source rent-or-buy algorithm gives a prize-collecting Steiner tree algorithm. We omit a discussion of the prize-collecting Steiner forest algorithm as it is obtained via an identical adaptation of the multi-source rent-or-buy algorithm.

We remark that the analyses for single-source rent-or-buy, connected facility location and prize-collecting Steiner tree rely on a property of the greedy online Steiner tree algorithm (Lemma 3.2), and the multi-source rent-or-buy algorithm depends on a property of the Berman-Coulston online Steiner forest algorithm (Lemma 4.1).

2 Basics of HST Embeddings

In this section, we present the necessary metric embedding definitions and results. Let (X, d) be a metric over a set of k terminals with the smallest distance $\min_{u,v \in X} d(u, v) = 1$. We are only interested in *expanding* embeddings, i.e. embeddings that do not shrink distances.

Definition 1 (Embeddings)**.** *A metric (X', d') is an* α-distortion embedding *of (X', d') with* distortion $\alpha \geq 1$ *if $X' \supseteq X$ and for all $u, v \in X$, we have $d(u, v) \leq d'(u, v) \leq \alpha d(u, v)$.*

In this paper, we will be concerned only with embeddings into hierarchically separated trees [4].

Definition 2 (HST Embeddings)**.** *A* hierarchically separated tree (HST) embedding *T of (X, d) is a rooted tree with height $\lceil \log(\max_{u,v \in X} d(u, v)) \rceil$ and edge lengths that are powers of 2 satisfying the following properties.*

1. *The leaves of T are exactly the terminals X.*
2. *The distance from any node to each of its children is the same.*
3. *The edge lengths decrease by a factor of 2 as one moves along a root-to-leaf path.*
4. *For an edge e in T, let $C_e \subseteq X$ be the subset of terminals that are separated by e from the root. We require that if e has length 2^{j-1}, then $d(u, v) < 2^j$ for any $u, v \in C_e$.*

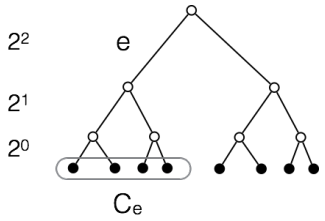

Figure 1: A binary HST embedding with 3 levels. The terminals X are shown in solid black. The edge e is a level-3 edge with length 2^2 and C_e is the set of terminals it separates from the root.

The last property will be crucial to our analyses. See Figure 1 for an example of a HST embedding.

A HST embedding T of a metric space (X, d) defines a hierarchical decomposition of (X, d) in the following way. Call an edge e of length 2^{j-1} a *level-j edge* of T and C_e a *level-j cut*. Denote by $E_j(T)$ the set of level-j edges and define $\mathcal{C}_j(T) = \{C_e : e \in E_j(T)\}$ to be the set of all level-j cuts. Now since the leaves of T are exactly X, the level-j cuts $\mathcal{C}_j(T)$ partitions X into subsets of diameter less than 2^j. The family of partitions $\mathcal{C}(T) = \{\mathcal{C}_j(T)\}_j$ is called the *hierarchical decomposition* of (X, d) defined by T. Note that there are $\lfloor \log(\max_{u,v \in X} d(u, v)) \rfloor$ levels and the level-0 cuts $\mathcal{C}_0(T)$ are just the terminal singletons.

In the following, we will use the notation $\delta(C)$ to denote the set of vertex pairs with exactly one endpoint in C, i.e. $\delta(C) = \{(u, v) : |\{u, v\} \cap C| = 1\}$. Note that $(u, v) \in \delta(C_e)$ if and only if e lies on the path between u and v in T.

We denote by $T(u, v)$ the distance between u, v in T. Fakcharoenphol et al. [8] showed that any metric can be embedded into HSTs with logarithmic expected distortion.

Theorem 2.1 ([8])**.** *For any metric (X, d), there exists a distribution $\mathcal{D}$ over HST embeddings T such that $E_{T \sim \mathcal{D}}[T(u, v)] \leq O(\log k) d(u, v)$ for all $u, v \in X$.*

The following corollary follows by standard arguments. We will apply it to the Steiner network and rent-or-buy problems.

Corollary 2.2. *For any network design instance with terminals X whose objective is a linear combination of edge lengths $d(u, v)$, there exists a distribution $\mathcal{D}$ over HST embeddings T of (X, d) such that $E_{T \sim \mathcal{D}}[\mathrm{OPT}(T)] \leq O(\log k)\,\mathrm{OPT}$, where OPT (and $\mathrm{OPT}(T)$ resp.) is the cost of the optimal solution on (X, d) (and T resp.).*

Again, we emphasize that no particular property of the distribution $\mathcal{D}$ is required. We only need the existence of a HST embedding T^* such that $\mathrm{OPT}(T^*) \leq O(\log k)\,\mathrm{OPT}$.

3 Warm up: Steiner Tree

We warm up by applying our approach to the online Steiner tree problem and give an alternative proof that the greedy algorithm is $O(\log k)$-competitive [14] in Section 3. Recall the problem statement. The algorithm is given a root terminal r at the beginning. Terminals i arrive online and the algorithm maintains a subgraph H connecting terminals to the root. The greedy algorithm is very natural — when terminal i arrives, it connects i to the nearest previously-arrived terminal.

Analysis. Let $X \subseteq V$ be the set of terminals that arrived and (X, d) be the submetric induced by X. The total cost of the greedy algorithm is $\sum_i a_i$ where a_i is the distance between terminal i and the nearest previously-arrived terminal. Our goal is to show that we can charge the cost of the greedy algorithm against the cost of the optimal solution on any HST embedding of the terminals.

Following the proof strategy outlined in the introduction, we first define cost shares for each terminal such that the total cost share accounts for the cost of the algorithm. We classify terminals according to a_i

— define class$(i) = j$ if $a_i \in [2^j, 2^{j+1})$ and $Z_j \subseteq X$ to be the set of class-j terminals. We define the cost share of each class-j terminal $i \in Z_j$ to be 2^{j+1}, i.e. i's cost share is a_i rounded up to the next power of 2. Thus, the total cost share is at least the cost of the algorithm.

Lemma 3.1. $\sum_i a_i \leq \sum_j 2^{j+1}|Z_j|$.

Before we proceed, we show that class-j terminals are at least 2^j-apart from each other.

Lemma 3.2. *For any terminals $i, i' \in Z_j$, we have that $d(i, i') \geq 2^j$.*

Proof. Suppose i arrived later than i'. We have $a_i \leq d(i, i')$ since i could have connected to i', and $a_i \geq 2^j$ by definition of Z_j. Thus, we get $d(i, i') \geq 2^j$. □

Next, we show that we can charge the cost shares against the cost of the optimal solution on any HST embedding of the terminals.

Lemma 3.3. $\sum_j 2^{j+1}|Z_j| \leq O(1)\,\mathrm{OPT}(T)$ *for any HST embedding T of (X, d).*

Proof. Let T be a HST embedding of (X, d). We begin by expressing $\mathrm{OPT}(T)$ in terms of cuts from the hierarchical decomposition $\mathcal{C}(T)$. Since the terminals are exactly the leaves of T, the unique feasible solution on T is the entire tree. In particular, we have

$$\mathrm{OPT}(T) = \sum_j 2^{j-1}|E_j(T)| = \sum_j 2^{j-1}|\mathcal{C}_j(T)|.$$

For each terminal $i \in Z_j$, we charge i's cost share to the level-j cut containing it. The total charge received by a level-j cut $C \in \mathcal{C}_j(T)$ is $2^{j+1}|Z_j \cap C|$ since we charge 2^{j+1} for each class-j terminal in C. Thus we have

$$\sum_j 2^{j+1}|Z_j| = \sum_j 2^{j+1} \sum_{C \in \mathcal{C}_j(T)} |Z_j \cap C|.$$

Each level-j cut $C \in \mathcal{C}_j(T)$ has diameter less than 2^j so Lemma 3.2 implies $|Z_j \cap C| \leq 1$. Therefore, we have $\sum_j 2^{j+1}|Z_j| \leq \sum_j 2^{j+1}|\mathcal{C}_j(T)| = 4\,\mathrm{OPT}(T)$. □

We are now ready to bound the competitive ratio of the greedy algorithm. Lemmas 3.1 and 3.3 imply that the cost of the greedy algorithm is at most $O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of (X, d). Furthermore, Corollary 2.2 implies that there exists a HST embedding T^* such that $\mathrm{OPT}(T^*) \leq O(\log k)\,\mathrm{OPT}$. Thus, the greedy algorithm is $O(\log k)$-competitive for the online Steiner tree problem.

4 Steiner Network

In this section, we consider the Steiner network problem and prove Theorem 1.1. Recall that in the online Steiner network problem, a terminal pair (s_i, t_i) with requirement R_i arrives at each online step and the algorithm maintains a multigraph H which contains R_i edge-disjoint (s_i, t_i)-paths for each arrived terminal pair. Let $R_{\max} = \max_i R_i$ be the maximum requirement. Note that the Steiner forest problem is a special case with $R_{\max} = 1$.

Intuition. Consider a "scaled" online Steiner forest instance in which all terminal pairs have the same requirement R. Since we can buy multiple copies of an edge, we can just run the $O(\log k)$-competitive Berman-Coulston algorithm and buy R copies of every edge it buys. The cost of the solution found is $O(\log k)\,\mathrm{OPT}$. This suggests the following approach: decompose the online Steiner network instance into $O(\log R_{\max})$ scaled Steiner forest subinstances and use the Berman-Coulston algorithm on each of them.

Agrawal et al. [1] applied this approach in the offline setting (using their offline Steiner forest algorithm) and obtained a $O(\log R_{\max})$-approximation, losing a factor of $O(\log R_{\max})$ over the approximation factor of their Steiner forest algorithm. There is a matching $\Omega(\log \min\{R_{\max}, k\})$ lower bound for this algorithm as

well. Intuitively, the $O(\log R_{\max})$ loss makes sense since each of the $O(\log R_{\max})$ subinstances are treated independently. Surprisingly, we show that in the online setting, this approach is $O(\log k)$-competitive, losing only a constant factor over the competitive ratio of the Berman-Coulston algorithm.

The analysis consists of two ingredients. The first is that for any online Steiner forest instance, the cost of the Berman-Coulston algorithm is at most a constant times the cost of the optimal solution on any HST embedding of the terminals (Lemma 4.2). The second is the following key idea: for a Steiner network instance on a tree, the total cost of the optimal solutions to the scaled Steiner forest subinstances is at most a constant times the cost of the optimal solution to the original Steiner network instance[3].

4.1 Steiner Forest

We now show that given an online Steiner forest instance, the Berman-Coulston algorithm (we call it Algorithm 1 henceforth) finds a solution whose cost is at most a constant times the optimal solution on any HST embedding of the terminals. We remark that this also gives a simpler analysis of the algorithm.

Algorithm. Algorithm 1 proceeds as follows. When a terminal pair (s_i, t_i) arrives, it classifies s_i and t_i based on their distance: $\text{class}(s_i) = \text{class}(t_i) = \lfloor \log d(s_i, t_i) \rfloor$. The algorithm then proceeds in *levels*, starting from level $j = 0$ up to class(i); at level j, it connects each terminal v of class at least j to s_i if $d(s_i, v) < 2^{j+1}$, or to t_i if $d(t_i, v) < 2^{j+1}$. We emphasize that we always consider distances according to the original metric d; we do not contract added edges.

Algorithm 1 Berman-Coulston Algorithm for Steiner Forest

```
H ← ∅
while request (s_i, t_i) arrives do
    Set class of s_i and t_i to be ⌊log d(s_i, t_i)⌋
    for level j = 0 to ⌊log d(s_i, t_i)⌋ do
        for v such that class(v) ≥ j and d(s_i, v) < 2^{j+1} do
            Add (s_i, v) to H
        end for
        for v such that class(v) ≥ j and d(t_i, v) < 2^{j+1} do
            Add (t_i, v) to H
        end for
    end for
end while
```

Analysis. Let X be the set of terminals that arrived and define X_j to be the set of terminals t with $\text{class}(t) \geq j$. Let A_j be the set of edges added in level j of some iteration. Note that $(u, v) \in A_j$ implies that $d(u, v) < 2^{j+1}$ and both u and v have class at least j.

We need the following lemma for our charging scheme.

Lemma 4.1. *For each class j, let $\mathcal{S}_j$ be a collection of disjoint subsets of X such that*

- *$\mathcal{S}_j$ covers X_j and*
- *for each $S \in \mathcal{S}_j$, we have $d(u, v) < 2^j$ for $u, v \in S$.*

Then we have $c(H) \leq \sum_j 2^{j+1} |\mathcal{S}_j|$.

Proof. We have that $c(H) \leq \sum_j 2^{j+1} |A_j|$ and so it suffices to prove that $|A_j| \leq |\mathcal{S}_j|$ for each j. Fix j and consider the following meta-graph: nodes correspond to $\mathcal{S}_j$; for each edge $(u, v) \in A_j$, there is a meta-edge

[3] In general graphs, we lose a factor of $O(\log R_{\max})$.

between the nodes corresponding to sets containing u and v. This is well-defined since $u, v \in X_j$ and $\mathcal{S}_j$ covers X_j. The meta-edges correspond to A_j and the number of nodes is $|\mathcal{S}_j|$. We will show that the meta-graph is acyclic and so $|A_j| \leq |\mathcal{S}_j|$, as desired.

Suppose, towards a contradiction, that there is a cycle in the meta-graph. Thus, there exists edges $(u_1, v_1), \ldots, (u_\ell, v_\ell) \in A_j$ and sets $S_1, \ldots, S_\ell \in \mathcal{S}_j$ such that $v_i, u_{i+1} \in S_{i+1}$ for each $i < \ell$ and $v_\ell, u_1 \in S_1$. Suppose that (u_ℓ, v_ℓ) was the last edge that was added by the algorithm. Since $(u_\ell, v_\ell) \in A_j$, it was added in level j of some iteration. Consider the point in time right before the algorithm added (u_ℓ, v_ℓ). For $1 \leq i \leq \ell$, the terminals $X_j \cap S_i$ that have arrived by this iteration are already connected by now since $\operatorname{diam}(S_i) < 2^j$ and they have class at least j. Therefore u_ℓ and v_ℓ are actually already connected at this time, and so the algorithm would not have added the edge (u_ℓ, v_ℓ). This gives the desired contradiction and so the meta-graph is acyclic, as desired. □

Lemma 4.2. *Define the function $f : 2^X \to \{0, 1\}$ as follows: $f(S) = 1$ if S separates a terminal pair. For any HST embedding T of (X, d), we have $c(H) \leq 4 \sum_j \sum_{C \in \mathcal{C}_j} 2^{j-1} f(C) = 4 \operatorname{OPT}(T)$.*

Proof. Let T be a HST embedding. Define $\mathcal{S}_j = \{C \in \mathcal{C}_j(T) : C \cap X_j \neq \emptyset\}$, i.e. $\mathcal{S}_j$ is the collection of level-j cuts that contain a terminal of class at least j. Since $\mathcal{S}_j$ satisfies the conditions of Lemma 4.1, we have that

$$c(H) \leq \sum_j 2^{j+1} |\mathcal{S}_j|.$$

Next, we lower bound $\operatorname{OPT}(T)$ in terms of its hierarchical decomposition $\mathcal{C}(T)$. For any level-j cut $C_e \in \mathcal{C}_j(T)$, if $(s_i, t_i) \in \delta(C_e)$ then e lies on the (s_i, t_i) path in T so the optimal solution on T buys e. Thus, we have that

$$\operatorname{OPT}(T) = \sum_j \sum_{C \in \mathcal{C}_j(T)} 2^{j-1} f(C).$$

Since $c(H) \leq \sum_j 2^{j+1} |\mathcal{S}_j|$ and $\mathcal{C}_j(T) \supseteq \mathcal{S}_j$, it suffices to prove that $f(C) = 1$ for $C \in \mathcal{S}_j$. Fix a cut $C \in \mathcal{S}_j$. By definition of $\mathcal{S}_j$, there exists a terminal $s_i \in C$ with $\operatorname{class}(s_i) \geq j$. Since $\operatorname{class}(s_i) \geq j$, we have that $d(s_i, t_i) \geq 2^j$ and so $t_i \notin C$ as the diameter of a level-j cut is less than 2^j. Thus, C separates the terminal pair (s_i, t_i) and so $f(C) = 1$, as desired. Now we have $\sum_{C \in \mathcal{C}_j(T)} f(C) \geq |\mathcal{S}_j|$ for each level j and this completes the proof of the lemma. □

Observe that this lemma together with Corollary 2.2 implies that Algorithm 1 is $O(\log k)$-competitive for the online Steiner forest problem. Lemma 4.2 says that the cost of Algorithm 1 is at most $O(1) \operatorname{OPT}(T)$ for any HST embedding T of (X, d). Furthermore, Corollary 2.2 implies that there exists a HST embedding T^* such that $\operatorname{OPT}(T^*) \leq O(\log k) \operatorname{OPT}$. Thus, Algorithm 1 is $O(\log k)$-competitive for the online Steiner forest problem.

4.2 From Steiner Forest to Steiner Network

We now state our Steiner network algorithm formally and show that it is $O(\log k)$-competitive.

Algorithm. For ease of exposition, we first assume that we are given the maximum requirement $R_{\max}$ and remove this assumption later on. We run $\lfloor \log R_{\max} \rfloor$ instantiations of Algorithm 1 (the Berman-Coulston Steiner forest algorithm); when we receive a terminal pair with requirement $R_i \in [2^\ell, 2^{\ell+1})$, we pass the pair to the ℓ-th instantiation and buy $2^{\ell+1}$ copies of each edge bought by that instantiation. The assumption that we are given $R_{\max}$ can be removed by starting a new instantiation of Algorithm 1 when we receive a terminal pair whose requirement is higher than all previous requirements.

Analysis. We maintain a feasible solution so it remains to bound its cost. Let X be the set of terminals that arrived and H_ℓ be the final subgraph of the ℓ-th instantiation of Algorithm 1. The cost of our solution is $\sum_\ell 2^{\ell+1} c(H_\ell)$. We now show that this is at most a constant times the cost of the optimal solution on any HST embedding of the terminals.

Lemma 4.3. $\sum_\ell 2^{\ell+1} c(H_\ell) \le O(1)\,\mathrm{OPT}(T)$ *for any HST embedding* T *of* (X, d).

Proof. Fix a HST embedding T of (X, d). Define the function $f : 2^X \to \mathbb{N}$ as follows: for $S \subseteq X$, we have $f(S) = \max_{i:(s_i,t_i)\in\delta(S)} R_i$. For any level-$j$ cut $C_e \in \mathcal{C}_j(T)$, if $(s_i, t_i) \in \delta(C_e)$ then e lies on the (s_i, t_i) path in T so the optimal solution on T buys $f(C_e)$ copies of e. Summing over all cuts of the hierarchical decomposition gives us

$$\mathrm{OPT}(T) = \sum_j \sum_{C\in\mathcal{C}_j(T)} 2^{j-1} f(C).$$

For each $0 \le \ell \le \lfloor \log R_{\max} \rfloor$, define the subset of terminal pairs $D_\ell = \{(s_i, t_i) : R_i \in [2^\ell, 2^{\ell+1})\}$ and the function $f_\ell : 2^X \to \{0, 1\}$ as follows: $f_\ell(S) = 1$ if $S \subseteq X$ separates a terminal pair of D_ℓ. Since H_ℓ is the output of Algorithm 1 when run on the subsequence of terminal pairs D_ℓ, Lemma 4.2 implies that

$$\begin{aligned}\sum_\ell 2^{\ell+1} c(H_\ell) &\le 4 \sum_\ell 2^{\ell+1} \left(\sum_j \sum_{C\in\mathcal{C}_j(T)} 2^{j-1} f_\ell(C) \right) \\ &= 4 \sum_j \sum_{C\in\mathcal{C}_j(T)} 2^{j-1} \left(\sum_\ell 2^{\ell+1} f_\ell(C) \right).\end{aligned}$$

Fix a cut C. Since each terminal pair $(s_i, t_i) \in D_\ell$ has requirement $R_i \in [2^\ell, 2^{\ell+1})$ and $f_\ell(C) = 1$ if C separates a terminal pair of D_ℓ, we have $\sum_\ell 2^{\ell+1} f_\ell(C) \le 4 \max_{i:(s_i,t_i)\in\delta(C)} R_i = 4 f(C)$. Thus, we have

$$4 \sum_j \sum_{C\in\mathcal{C}_j(T)} 2^{j-1} \left(\sum_\ell 2^{\ell+1} f_\ell(C) \right) \le 16 \sum_j \sum_{C\in\mathcal{C}_j(T)} 2^{j-1} f(C) = 16\,\mathrm{OPT}(T).$$

□

By Corollary 2.2, there exists a HST embedding T^* such that $\mathrm{OPT}(T^*) \le O(\log k)\,\mathrm{OPT}$. So Lemma 4.3 implies that the cost of our algorithm is at most $O(\log k)\,\mathrm{OPT}$. Thus, our algorithm is $O(\log k)$-competitive for Steiner network and this proves Theorem 1.1.[4]

5 Rent-or-Buy

Next, we consider the rent-or-buy problem and prove Theorem 1.2. We illustrate the key ideas by proving the single-source version of the theorem in Section 5.1. Then we extend these ideas to the multi-source setting in Section 5.2.

5.1 Single-Source Rent-or-Buy

We prove the following theorem.

Theorem 5.1. *There exists a* $O(\log k)$*-competitive algorithm for the online single-source rent-or-buy problem.*

[4]We remark that one can apply the usual dual-fitting analysis, but the upper bound it gives is proportional to the number of relevant distance scales. This is $O(\log(kR_{\max}))$ for this problem and can deteriorate to $O(k)$ since $R_{\max}$ can be as large as 2^k.

We recall the problem statement. The algorithm is initially given a root terminal r and a parameter $M \geq 0$. The algorithm maintains a subgraph H of bought edges. When a terminal i arrives, the algorithm buys zero or more edges and rents edges Q_i such that $H \cup Q_i$ connects i and r. The total cost of the algorithm is $Mc(H) + \sum_i c(Q_i)$. We call $Mc(H)$ the *buy cost* and $c(Q_i)$ the *rent cost* of terminal i.

Intuition. Our algorithm maintains that the subgraph H of bought edges is a connected subgraph containing the root; essentially it either rents or buys the shortest edge from the current terminal to H. Roughly speaking, our algorithm buys if there are M terminals nearby with sufficiently large rent costs, and rents otherwise. This allows us to charge the buy cost to the total rent cost and we then charge the total rent cost to edges of any HST embedding T by showing that there are few terminals with large rent cost in a small neighborhood.

Algorithm. Algorithm 2 maintains a set of *buy terminals* Z (Z includes r) which are connected by H to the root; all other terminals are called *rent terminals* and denoted by R. When a terminal i arrives, let $z \in Z$ be the closest buy terminal and define $a_i = d(i, z)$. As in the previous section, we define $\text{class}(i) = j$ if $a_i \in [2^j, 2^{j+1})$. The algorithm considers the set of class-j rent terminals R_j of distance less than 2^{j-1} to i. Call this the *witness set* of i and denote it by $W(i)$. The algorithm buys the edge (i, z) if $|W(i)| \geq M$ and rents it otherwise. Terminal i becomes a buy terminal in the former case and a rent terminal in the latter case.

Algorithm 2 Algorithm for Online Single-Source Rent-or-Buy

$Z \leftarrow \{r\}; R_j \leftarrow \emptyset; H \leftarrow \emptyset$
while terminal i arrives **do**
  Let z be closest terminal in Z to i and set $j = \lfloor \log d(i, z) \rfloor$
  $W(i) \leftarrow \{i' \in R_j : d(i, i') < 2^{j-1}\}$
  if $|W(i)| \geq M$ **then**
    Add i to Z and buy (i, z), i.e. $H \leftarrow H \cup \{(i, z)\}$
  else
    Add i to R_j and rent (i, z), i.e. $Q_i \leftarrow \{(i, z)\}$
  end if
end while

Analysis. Let X be the set of all terminals. For each buy terminal $z \in Z$, the algorithm incurs a buy cost of Ma_z; for each rent terminal $i \in R$, the algorithm incurs a rent cost of a_i. Thus, the total cost of the algorithm is $\sum_{z \in Z} Ma_z + \sum_{i \in R} a_i$. Our goal in the analysis is to show that this is at most a constant times $\text{OPT}(T)$ for any HST embedding T of (X, d).

First, we define cost shares for each terminal. For each rent terminal $i \in R_j$, we define i's cost share to be 2^{j+1}, i.e. a_i rounded up to the next power of 2. The total cost share is $\sum_j 2^{j+1}|R_j|$. We now show that the cost of the algorithm is at most twice the total cost share.

Lemma 5.2. $\sum_{z \in Z} Ma_z + \sum_{i \in R} a_i \leq 2 \sum_j 2^{j+1}|R_j|$.

Proof. Let $Z_j \subseteq Z$ be the set of class-j buy terminals. We have $\sum_{i \in R} a_i \leq \sum_j 2^{j+1}|R_j|$ and $\sum_{z \in Z} Ma_z \leq \sum_j M 2^{j+1}|Z_j|$. We now show that $M|Z_j| \leq |R_j|$ for each class j.

Fix a class j. The witness set $W(z)$ of $z \in Z_j$ satisfies the following properties: (1) $|W(z)| \geq M$; (2) $W(z) \subseteq R_j$; (3) $d(i, z) < 2^{j-1}$ for $i \in W(z)$. The first implies that $M|Z_j| \leq \sum_{z \in Z_j} |W(z)|$. We claim that $d(z, z') \geq 2^j$ for $z, z' \in Z_j$. This completes the proof since together with the second and third properties, we have that the witness sets of Z_j are disjoint subsets of R_j and so $\sum_{z \in Z_j} |W(z)| \leq |R_j|$.

Now we prove the claim. Observe that H is the subgraph produced by the greedy online Steiner tree algorithm if it were run on the subsequence of buy terminals Z and for each $z \in Z$, we have that a_z is

exactly the distance from z to the nearest previously-arrived buy terminal. Thus, we can apply Lemma 3.2 and get that $d(z, z') \geq 2^j$ for any $z, z' \in Z_j$, proving the claim. Putting all of the above together, we get $\sum_{z \in Z} M a_z \leq \sum_j 2^{j+1}|R_j|$. This finishes the proof of the lemma. □

Next, we show that the total cost share is at most a constant times the cost of the optimal solution on any HST embedding of the terminals.

Lemma 5.3. *$\sum_j 2^{j+1}|R_j| \leq O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of (X, d).*

Proof. Let T be a HST embedding of (X, d). We begin by lower bounding $\mathrm{OPT}(T)$ in terms of T's cuts. Fix a level-j cut $C_e \in \mathcal{C}_j(T)$. If $r \notin C_e$, then e lies on the path between i and r for each terminal $i \in C_e$. Since $e \in E_j(T)$ and has length 2^{j-1}, the optimal solution on T either rents e for each terminal in C at a cost of $2^{j-1}|C|$ or buys it at a cost of $2^{j-1}M$. Summing over all cuts of the hierarchical decomposition gives us the following lower bound

$$\mathrm{OPT}(T) \geq \sum_j \sum_{C \in \mathcal{C}_j(T): r \notin C} 2^{j-1} \min\{M, |C|\}.$$

For each rent terminal $i \in R_j$, we charge 2^{j+1} to the unique level-$(j-1)$ cut[5] containing i. Each level-j cut $C \in \mathcal{C}_j(T)$ is charged $2^{j+2}|R_{j+1} \cap C|$ since it receives a charge of 2^{j+2} for each class-$(j+1)$ rent terminal in it. So overall, we have

$$\sum_j 2^{j+1}|R_j| = \sum_j \sum_{C \in \mathcal{C}_j(T)} 2^{j+2}|R_{j+1} \cap C|.$$

It remains to prove the following claim: for each level-j cut $C \in \mathcal{C}_j(T)$, we have $|R_{j+1} \cap C| = 0$ if $r \in C$ and $|R_{j+1} \cap C| \leq \min\{M, |C|\}$ if $r \notin C$. Suppose $r \in C$. Since $a_i < d(i, r) < 2^j$ for all $i \in C$, there cannot be any class-$(j+1)$ terminal in C and so $R_{j+1} \cap C = \emptyset$. Now consider the case $r \notin C$. Since $|R_{j+1} \cap C| \leq |C|$, it suffices to prove that $|R_{j+1} \cap C| \leq M$. Suppose, towards a contradiction, that $|R_{j+1} \cap C| > M$ and let i be the last-arriving terminal of $R_{j+1} \cap C$. The terminals of $R_{j+1} \cap C \setminus \{i\}$ arrive before i, are each of distance less than 2^j from i (diameter of C is less than 2^j) and of the same class as i, so they are part of i's witness set $W(i)$. Since $|R_{j+1} \cap C| > M$, we have that $|W(i)| \geq |R_{j+1} \cap C \setminus \{i\}| \geq M$. Thus, i would have been a buy terminal but this contradicts the assumption that $i \in R_{j+1}$. Therefore, we have $|R_{j+1} \cap C| \leq M$ as desired. This completes the proof of the claim and so we get $\sum_j 2^{j+1}|R_j| \leq 8\,\mathrm{OPT}(T)$. □

Now we put all of the above together. Lemmas 5.2 and 5.3 imply that the cost of Algorithm 2 is at most $O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of (X, d). Furthermore, by Corollary 2.2, there exists a HST embedding T^* such that $\mathrm{OPT}(T^*) \leq O(\log k)\,\mathrm{OPT}$. Thus the algorithm is $O(\log k)$-competitive for single-source rent-or-buy and this proves Theorem 5.1.

5.2 Multi-Source Rent-or-Buy

We move on to the multi-source setting. Recall the problem statement. The algorithm is given a parameter $M \geq 0$ initially, and it maintains a subgraph H of bought edges online. When a terminal pair (s_i, t_i) arrives, the algorithm buys zero or more edges and rents edges Q_i such that $H \cup Q_i$ connects s_i and t_i. The total cost of the algorithm is $Mc(H) + \sum_i c(Q_i)$. We call $Mc(H)$ the *buy cost* and $c(Q_i)$ the *rent cost* of terminal pair (s_i, t_i).

Intuition. The high-level idea is similar to the single-source rent-or-buy algorithm in Section 5.1. For each terminal pair (s_i, t_i), our algorithm will either rent the edge (s_i, t_i) or buy edges such that s_i and t_i are connected in H. Roughly speaking, the algorithm buys edges if there are at least M endpoints of terminal pairs with sufficiently large rent costs near each of s_i and t_i, and rents the edge (s_i, t_i) otherwise. The main difference with the single-source case is that H is a Steiner forest connecting a subset of the terminal pairs. Thus, we use the Berman-Coulston algorithm for online Steiner forest [6] (described in Section 4.1) to determine H.

[5] We can extend T with an additional level of terminal singletons to accomodate charging against level $j = -1$. This only increases $\mathrm{OPT}(T)$ by at most a constant.

Algorithm. For each terminal pair (s_i, t_i), we define $a_i = d(s_i, t_i)$ and classify s_i and t_i based on a_i: $\text{class}(s_i) = \text{class}(t_i) = j$ if $d(s_i, t_i) \in [2^j, 2^{j+1})$. As terminal pairs arrive online, the algorithm designates some of the terminals as *rent terminals*. When (s_i, t_i) arrives, the algorithm considers the sets $W(s_i)$ and $W(t_i)$ of class-j rent terminals that are of distance less than 2^{j-2} to s_i and t_i, respectively. We call these the *witness sets* of s_i and t_i. If $|W(s_i)| \geq M$ and $|W(t_i)| \geq M$, the algorithm buys edges such that s_i and t_i are connected in H; otherwise, it rents the edge (s_i, t_i) and designates exactly one of s_i or t_i to be a rent terminal. We say that (s_i, t_i) is a *buy pair* if the edge is bought and a *rent pair* otherwise.

In the following, denote by R_j the set of class-j rent terminals that have arrived so far.

Algorithm 3 Algorithm for Online Multi-Source Rent-or-Buy

$H \leftarrow \emptyset; R_j \leftarrow \emptyset$
while request (s_i, t_i) arrives **do**
  Set $j = \text{class}(s_i) = \text{class}(t_i) = \lfloor \log d(s_i, t_i) \rfloor$
  $W(s_i) \leftarrow \{v \in R_j : d(s_i, v) < 2^{j-2}\}$
  $W(t_i) \leftarrow \{v \in R_j : d(t_i, v) < 2^{j-2}\}$
  if $|W(s_i)| < M$ **then**
    Rent (s_i, t_i), i.e. $Q_i \leftarrow \{(s_i, t_i)\}$
    Add s_i to R_j
  else if $|W(t_i)| < M$ **then**
    Rent (s_i, t_i), i.e. $Q_i \leftarrow \{(s_i, t_i)\}$
    Add t_i to R_j.
  else
    Pass (s_i, t_i) to Algorithm 1 and buy each edge that it buys; add bought edges to H
  end if
end while

Analysis. Let X be the set of terminals that arrived, R be the set of all rent terminals and Z be the set of terminals that were passed to Algorithm 1 (i.e. Z is the set of endpoints of buy pairs). We call Z the set of *buy terminals*. Note that R contains exactly one endpoint of each rent pair; we rename the terminals so that if (s_i, t_i) is a rent pair, then $t_i \in R$. The cost of the algorithm is $Mc(H) + \sum_{t_i \in R} d(s_i, t_i)$. The analysis proceeds by charging the cost of the algorithm against the cost of the optimal solution on any HST embedding of the terminals.

We first define cost shares for each terminal. For each $t_i \in R_j$, we define t_i's cost share to be 2^{j+1}. The total cost share is $\sum_j 2^{j+1}|R_j|$. We now show that the cost of the algorithm is at most twice the total cost share.

Lemma 5.4. $Mc(H) + \sum_{t_i \in R} d(s_i, t_i) \leq 2 \sum_j 2^{j+1}|R_j|$.

Proof. Since t_i's cost share is at least $d(s_i, t_i)$, we have that $\sum_{t_i \in R} d(s_i, t_i) \leq \sum_j 2^{j+1}|R_j|$. Next, we bound $Mc(H)$. Let $Z_j \subseteq Z$ be the set of buy terminals with class at least j. Define $Z'_j \subseteq Z_j$ to be the maximal subset of Z_j such that $d(u, v) \geq 2^{j-1}$ for all $u, v \in Z'_j$. We will show that $Mc(H) \leq \sum_j 2^{j+1}|R_j|$ in two steps: first we prove that $Mc(H) \leq \sum_j 2^{j+1} M|Z'_j|$ and then prove that $M|Z'_j| \leq |R_j|$ for each j.

We partition Z_j as follows: assign each $v \in Z_j$ to the closest terminal in Z'_j, breaking ties arbitrarily, and define S_u to be the set of terminals assigned to u. Observe that the diameter of S_u is less than 2^j for all $u \in Z'_j$. Define $\mathcal{S}_j = \{S_u\}_{u \in Z'_j}$. Since $\mathcal{S}_j$ satisfies the conditions of Lemma 4.1, we get that $Mc(H) \leq \sum_j 2^{j+1} M|\mathcal{S}_j| = \sum_j 2^{j+1} M|Z'_j|$.

Next we show that $M|Z'_j| \leq 2^{j+1}|R_j|$ for each j. Fix j. The witness set $W(z)$ of $z \in Z'_j$ satisfies the following properties: (1) $|W(z)| \geq M$; (2) $W(z) \subseteq R_j$; (3) $d(i, z) < 2^{j-2}$ for $i \in W(z)$. The first implies that $M|Z_j| \leq \sum_{z \in Z_j} |W(z)|$. We have $d(z, z') \geq 2^{j-1}$ for $z, z' \in Z'_j$ by definition, so with the second and

third properties, we get that the witness sets of Z'_j are disjoint subsets of R_j, and so $\sum_{z \in Z'_j} |W(z)| \leq |R_j|$. Putting all of the above together, we have $Mc(H) \leq \sum_j 2^{j+1}|R_j|$. □

Next, we show that we can charge the cost shares against the optimal solution on any HST embedding T.

Lemma 5.5. $\sum_j 2^{j+1}|R_j| \leq O(1)\,\mathrm{OPT}(T)$ *for all HST embeddings* T *of* (X, d).

Proof. Let T be a HST embedding of (X, d). We begin by expressing $\mathrm{OPT}(T)$ in terms of T's cuts. Define $D(S) = \{(s_i, t_i) : (s_i, t_i) \in \delta(S)\}$ for each vertex subset $S \subseteq X$. Consider a level-j cut $C_e \in \mathcal{C}_j(T)$. By definition, $e \in E_j(T)$ and has length 2^{j-1}. If $(s_i, t_i) \in \delta(C_e)$, then e lies on the (s_i, t_i) path in T. Thus, the optimal solution on T either rents e for each terminal pair in $D(C_e)$ at a cost of $2^{j-1}|D(C_e)|$ or buys it at a cost of $2^{j-1}M$ and so

$$\mathrm{OPT}(T) = \sum_j \sum_{C \in \mathcal{C}_j(T)} 2^{j-1} \cdot \min\{M, |D(C)|\}.$$

For each rent terminal $t_i \in R_j$, we charge its cost share 2^{j+1} to the level-$(j-2)$ cut $C \in \mathcal{C}_{j-2}(T)$ containing i. Each level-j cut $C \in \mathcal{C}_j(T)$ receives a charge of $2^{j+3}|R_{j+2} \cap C|$ and so

$$\sum_j 2^{j+1}|R_j| = \sum_j 2^{j+3} \sum_{C \in \mathcal{C}_j(T)} |R_{j+2} \cap C|.$$

It remains to prove the following claim: for each level-j cut $C \in \mathcal{C}_j(T)$, we have $|R_{j+2} \cap C| \leq \min\{M, |D(C)|\}$. Since C has diameter less than 2^j and $d(s_i, t_i) \geq 2^{j+2}$ for each $t_i \in R_{j+2}$, we have $|R_{j+2} \cap C| \leq |D(C)|$. So now we prove that $|R_{j+2} \cap C| \leq M$. Suppose, towards a contradiction, that $|R_{j+2} \cap C| > M$ and let t_i be the last-arriving terminal of $R_{j+2} \cap C$. The terminals of $R_{j+2} \cap C \setminus \{t_i\}$ arrive before t_i, are each of distance less than 2^j from t_i (diameter of C is less than 2^j) and of the same class as t_i, so they are part of t_i's witness set $W(t_i)$. Since $|R_{j+2} \cap C| > M$, we have that $|W(t_i)| \geq |R_{j+2} \cap C \setminus \{t_i\}| \geq M$. Thus, t_i would have been a buy terminal but this contradicts the assumption that $t_i \in R_{j+2}$. Therefore, we have $|R_{j+2} \cap C| \leq M$ as desired. This completes the proof of the claim and so $\sum_j 2^{j+1}|R_j| \leq 16\,\mathrm{OPT}(T)$. □

Now, Lemmas 5.4 and 5.5 imply that the cost of Algorithm 3 is at most $O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of (X, d). Furthermore, by Corollary 2.2, there exists a HST embedding T^* such that $\mathrm{OPT}(T^*) \leq O(\log k)\,\mathrm{OPT}$. Thus the algorithm is $O(\log k)$-competitive for the online rent-or-buy problem and this proves Theorem 1.2.

6 Connected Facility Location

In this section, we consider the connected facility location problem and prove Theorem 1.3. We recall the problem statement. At the beginning, the algorithm is given a parameter $M \geq 0$, a set of *facilities* $F \subseteq V$ and facility opening costs f_x for each facility $x \in F$. There is a designated root facility $r \in F$ with zero opening cost. The algorithm maintains a set of open facilities F' and a subgraph H connecting F' and r. When a client i arrives, the algorithm may open a new facility, and then assigns i to some open facility $\sigma(i) \in F'$. The cost of the algorithm is $\sum_{x \in F'} f_x + Mc(H) + \sum_i d(i, \sigma(i))$. We call $Mc(H)$ the *Steiner cost* and $d(i, \sigma(i))$ the *assignment cost* of client i. The special case when $M = 0$ (i.e. open facilities need not be connected) is called the online facility location problem.

Intuition. The connected facility location problem shares a similar cost structure with single-source rent-or-buy. Here, the "buy cost" consists of the facility opening cost and the Steiner cost, the "rent cost" is the assignment cost. At a high-level, we use a similar strategy — we only open a new facility if we can pay for the facility opening cost and the additional Steiner cost using the assignment costs of nearby clients. In order to do this, we use an online facility location algorithm together with an adaptation of our single-source

rent-or-buy algorithm. The former tells us if we can pay for the opening cost and the latter tells us if we can pay for the additional Steiner cost incurred in connecting the newly open facility to r.

In the analysis, we will not be able to show that the cost of the algorithm is at most $O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of the clients because there is a lower bound of $\Omega(\frac{\log k}{\log\log k})$ for the facility location problem even on HSTs [10]. In a sense, routing is easy on trees but choosing a good set of facilities to open online remains difficult. Instead, we will use the fact that the connected facility location instance induces a rent-or-buy instance and a facility location instance on the same metric, and that there is a $O(\log k)$-competitive facility location algorithm [9][6]. Our analysis charges part of the cost to the cost of the facility location algorithm and the rest to the optimal rent-or-buy solution.

Algorithm. Since the root facility has zero opening cost, our algorithm starts by opening the root facility. We run Fotakis's $O(\log k)$-competitive algorithm [9] in parallel; call it OFL-ALG. We denote its open facilities by $\hat{F}$ and its assignments by $\hat{\sigma}$, and call them *virtual facilities* and *virtual assignments*, respectively. Algorithm 4 will only open a facility if it was already opened by OFL-ALG. i.e. its open facilities F' is a subset of the virtual facilities $\hat{F}$. (We assume w.l.o.g. that OFL-ALG opens the root facility as well.)

When a client i arrives, let $x \in F'$ be the nearest open facility. Define $a_i = d(i,x)$ and $\mathrm{class}(i) = j$ if $a_i \in [2^j, 2^{j+1})$. We first pass the client to OFL-ALG, which may open a new virtual facility and then it assigns i to the nearest virtual facility $\hat{\sigma}(i)$. We will either assign i to $\hat{\sigma}(i)$ or the nearest open facility x. (Note that $F' \subseteq \hat{F}$ so i is closer to $\hat{\sigma}(i)$ than x.) If x is not much further than $\hat{\sigma}(i)$ — in particular, if $d(i,x) \le 4d(i,\hat{\sigma}(i))$ — then we assign i to x since we can later charge the assignment cost $d(i,x)$ to OFL-ALG's assignment cost $d(i,\hat{\sigma}(i))$. In this case, we call i a *virtual client*. Otherwise, we consider opening $\hat{\sigma}(i)$ and assigning i to it. We define the *witness set* $W(i)$ to be the class-j clients that are of distance at most 2^{j-2} from i. If there are at least M witnesses, then we open $\hat{\sigma}(i)$, assign i to $\hat{\sigma}(i)$ and connect $\hat{\sigma}(i)$ to the root via x (i.e. we add the edge $(\hat{\sigma}(i), x)$ to H); we call i a *buy client* in this case. Otherwise, then we simply assign i to x and call i a *rent client*.

In the description of Algorithm 4 below, we use Q_j, Z_j, R_j to keep track of the class-j virtual clients, buy clients and rent clients, resp.

Algorithm 4 Algorithm for Online Connected Facility Location

Initialize $F' \leftarrow \{r\}; H \leftarrow \emptyset; Q_j \leftarrow \emptyset; Z_j \leftarrow \emptyset; R_j \leftarrow \emptyset;$
while client i arrives **do**
  Pass i to OFL-ALG and update virtual solution $\hat{F}, \hat{\sigma}$
  Let $x \in F'$ be nearest open facility to i and set $j = \lfloor \log d(i,x) \rfloor$
  $W(i) \leftarrow \{i' \in R_j : d(i',i) < 2^{j-2}\}$
  if $a_i \le 4d(i,\hat{\sigma}(i))$ **then**
    Assign i to x and add i to Q_j
  else if $|W(i)| \ge M$ **then**
    Open $\hat{\sigma}(i)$ and add $\hat{\sigma}(i)$ to F'
    Add edge $(\hat{\sigma}(i), x)$ to H
    Assign i to $\hat{\sigma}(i)$ and add i to Z_j
  else
    Assign i to x and add i to R_j
  end if
end while

Analysis. Let X be the set of clients and the root, and k be the number of clients. Consider a facility location instance over metric (V,d) with the same facilities F and clients X as well as a rent-or-buy instance over metric (V,d) with terminals X and root r. Let $\mathrm{OPT}_{\mathrm{FL}}$ and $\mathrm{OPT}_{\mathrm{ROB}}$ be the cost of the respective

[6] We can also apply our techniques to show that the algorithm of [9] is $O(\log k)$-competitive but we omit the discussion as the analysis differs greatly from the analyses in this paper.

optimal solutions. A feasible solution to the connected facility location instance is feasible for both the facility location and rent-or-buy instances so we have the following lemma.

Lemma 6.1. $\mathrm{OPT}_{\mathrm{FL}} \leq \mathrm{OPT}$ *and* $\mathrm{OPT}_{\mathrm{ROB}} \leq \mathrm{OPT}$.

Let Q, Z, R be the set of virtual, buy, and rent clients, respectively. Note that a virtual or rent client $i \in Q \cup R$ has assignment cost $d(i, \sigma(i)) = a_i$ and a buy client $i \in Z$ has assignment cost $d(i, \sigma(i)) = d(i, \hat{\sigma}(i))$. Thus the cost of the algorithm is

$$\sum_{x \in F'} f_x + Mc(H) + \sum_{i \in Q \cup R} a_i + \sum_{i \in Z} d(i, \hat{\sigma}(i)).$$

We first charge the opening cost as well as the assignment cost of virtual and buy clients to the cost of the virtual solution. Then we charge the Steiner cost and the assignment cost of rent clients to $\mathrm{OPT}_{\mathrm{ROB}}$.

Lemma 6.2. $\sum_{x \in F'} f_x + \sum_{i \in Q} a_i + \sum_{i \in Z} d(i, \hat{\sigma}(i)) \leq O(\log k)\,\mathrm{OPT}$.

Proof. Our algorithm only opens a facility if it was already opened by OFL-ALG so $\sum_{x \in F'} f_x \leq \sum_{x \in \hat{F}} f_x$. Furthermore, the assignment cost of a virtual client $i \in Q$ is $a_i \leq 4d(i, \hat{\sigma}(i))$, and for each buy client $i \in Z$ we have $d(i, \sigma(i)) = d(i, \hat{\sigma}(i))$. Therefore, we have

$$\sum_{x \in F'} f_x + \sum_{i \in Q} a_i + \sum_{i \in Z} d(i, \hat{\sigma}(i)) \leq \sum_{x \in \hat{F}} f_x + 4 \sum_{i} d(i, \hat{\sigma}(i)).$$

Since OFL-ALG is a $O(\log k)$-competitive algorithm for facility location, the cost of its virtual solution is $\sum_{x \in \hat{F}} f_x + \sum_i d(i, \hat{\sigma}(i)) \leq O(\log k)\,\mathrm{OPT}_{\mathrm{FL}}$. The lemma now follows from Lemma 6.1. □

Next, we show that the Steiner cost and the assignment cost of rent clients $Mc(H) + \sum_{i \in R} a_i$ is at most a constant times the cost of the optimal rent-or-buy solution on any HST embedding of the terminals. This part of the analysis is analogous to that in Section 5.1. For each rent client $i \in R_j$, we define i's cost share to be 2^{j+1}. We now show that the Steiner cost and the assignment cost of rent clients is at most thrice the total cost share $\sum_j 2^{j+1}|R_j|$.

We will need the following lemma which says that class-j buy clients are at least 2^{j-1}-apart from each other.

Lemma 6.3. $d(z, z') \geq 2^{j-1}$ *for* $z, z' \in Z_j$.

Proof. By triangle inequality, we have $d(z, z') \geq d(z, \sigma(z')) - d(z', \sigma(z'))$. Suppose z arrived after z'. Since a_z is defined to be the distance from z to the nearest open facility when z arrived and the facility $\sigma(z')$ that z' was assigned to is open at this time, we have $d(z, \sigma(z')) \geq a_z$. Moreover, z' is a buy client and so we assigned it to a facility $\sigma(z')$ such that $d(z', \sigma(z')) < \frac{1}{4} a_{z'}$. We now have that $d(z, z') \geq a_z - \frac{1}{4} a_{z'}$. Since z, z' are of class j, we have $a_z \geq 2^j$ and $a_{z'} < 2^{j+1}$. Thus, $d(z, z') \geq 2^{j-1}$. □

Lemma 6.4. $Mc(H) + \sum_{i \in R} a_i \leq 3 \sum_j 2^{j+1}|R_j|$.

Proof. Since the cost share of a rent terminal $i \in R$ is at least a_i, it suffices to show that $Mc(H) \leq 2 \sum_{i \in R} a_i$. We will prove this using the following two claims.

Claim 6.1. $c(H) \leq \sum_{z \in Z} 2a_z$.

Proof. We proceed by charging the increase in $c(H)$ due to an open facility to the buy client that opened it. For each open facility $y \in F'$, define $z(y) \in Z$ to be the buy client that opened it and $x(y) \in F'$ to be the open facility that y was connected to when it was opened. So $c(H) = \sum_{y \in F'} d(y, x(y))$. We now show that $d(y, x(y)) \leq 2a_{z(y)}$ for each $x \in F'$.

Fix an open facility $y \in F'$. For brevity, we write z, and x, in place of $z(y)$, and $x(y)$, respectively. By triangle inequality, $d(y, x) \leq d(y, z) + d(z, x)$ so now we bound the right-hand side in terms of a_z. By

definition of the algorithm, at the time when z just arrived, before it opened y, we have that x was the nearest open facility and y was the nearest virtual facility. So $a_z = d(z,x)$ and $y = \hat{\sigma}(z)$. Furthermore, $d(z,\hat{\sigma}(z)) < \frac{1}{4}a_z$ since z is not a virtual client. Thus, $d(y,x) \le d(y,z) + d(z,x) \le 2a_z$, as desired.

Each open facility was opened by a distinct buy client so $c(H) = \sum_{y \in F'} d(y, x(y)) \le \sum_{z \in Z} 2a_z$. □

Claim 6.2. $\sum_{z \in Z} M a_z \le \sum_j 2^{j+1}|R_j|$.

Proof. We have $\sum_{z \in Z} M a_z \le \sum_j 2^{j+1} M |Z_j|$. We now prove that $M|Z_j| \le |R_j|$ for each class j. The witness set $W(z)$ of $z \in Z_j$ satisfies the following properties: (1) $|W(z)| \ge M$; (2) $W(z) \subseteq R_j$; (3) $d(i,z) < 2^{j-2}$ for $i \in W(z)$. The first implies that $M|Z_j| \le \sum_{z \in Z_j} |W(z)|$. Lemma 6.3 together with the second and third properties imply that the witness sets of Z_j are disjoint subsets of R_j and so $\sum_{z \in Z_j} |W(z)| \le |R_j|$. These two inequalities imply that $\sum_{z \in Z} M a_z \le \sum_j 2^{j+1}|R_j|$. □

Combining these claims, we have $Mc(H) \le \sum_{z \in Z} 2Ma_z \le 2\sum_j 2^{j+1}|R_j|$. Therefore, $Mc(H) + \sum_{i \in R} a_i \le 3\sum_j 2^{j+1}|R_j|$. □

Now we show that the total cost share is at most a constant times the cost of the optimal rent-or-buy solution on any HST embedding of the terminals.

Lemma 6.5. $\sum_j 2^{j+1}|R_j| \le O(1)\,\mathrm{OPT}_{\mathrm{ROB}}(T)$ *for any HST embedding* T *of* (X,d).

Proof. Let T be a HST embedding. We charge the cost share of class-j clients to the level-$(j-2)$ cuts. So overall, we have $\sum_j 2^{j+1}|R_j| = \sum_j \sum_{C \in \mathcal{C}_j(T)} 2^{j+3}|R_{j+2} \cap C|$. The rest of the proof proceeds as in the proof of Lemma 5.3. □

Now we have all the required ingredients to bound the competitive ratio of Algorithm 4. By Corollary 2.2 and Lemma 6.1, there exists a HST embedding T^* such that $\mathrm{OPT}_{\mathrm{ROB}}(T^*) \le O(\log k)\,\mathrm{OPT}_{\mathrm{ROB}} \le O(\log k)\,\mathrm{OPT}$. So Lemmas 6.4 and 6.5 imply that the Steiner cost and the assignment cost of rent clients is $Mc(H) + \sum_{i \in R} a_i \le O(\log k)\,\mathrm{OPT}$. Lemma 6.2 says that the remainder of the algorithm's cost is at most $O(\log k)\,\mathrm{OPT}$ as well. Thus it is $O(\log k)$-competitive for connected facility location and this proves Theorem 1.3.

7 Prize-Collecting Steiner Tree

In this section, we give a simple algorithm and analysis for the prize-collecting Steiner tree problem and prove the following theorem.

Theorem 7.1 ([17])**.** *There is a deterministic $O(\log k)$-competitive algorithm for the online prize-collecting Steiner tree problem.*

We recall the problem statement. The algorithm is given a root terminal r initially and maintains a subgraph H online. At each online step, a terminal i with *penalty* π_i arrive and the algorithm can either pay the penalty or augment H such that H connects i to the root. We say that i is a *penalty terminal* if the algorithm chose to pay i's penalty and denote by P the set of penalty terminals. The total cost of the algorithm is $c(H) + \sum_{i \in P} \pi_i$.

Algorithm. Algorithm 5 maintains a set of *buy terminals* Z (Z includes r) which it connects to the root. It also associates to each terminal i a cost share ρ_i. When a terminal i with penalty π_i arrives, let $z \in Z$ be the closest buy terminal. We define $a_i = d(i,z)$, and define $\mathrm{class}(i) = j$ if $a_i \in [2^j, 2^{j+1})$. The algorithm initializes i's cost share ρ_i to 0 and raises ρ_i until either the total cost share of class-j terminals within a radius of 2^{j-1} of i (including i) is at least 2^{j+1} (sufficient to pay for the edge (i,z)) or $\rho_i = \pi_i$. In the first case, we add the edge (i,z) to H; in the second case, we pay the penalty. We denote by X_j the set of class-j terminals and $W(i)$ the set of class-j terminals within a radius of 2^{j-1} of i (we call $W(i)$ the *witness set* of i).

Algorithm 5 Algorithm for Online Prize-Collecting Steiner Tree

$Z \leftarrow \{r\}; H \leftarrow \emptyset$; $X_j \leftarrow \emptyset$ for all j
while terminal i with penalty π_i arrives **do**
  Let z be closest terminal in Z to i, set $j = \lfloor \log d(i,z) \rfloor$ and add i to X_j
  $W(i) \leftarrow \{i' \in X_j : d(i,i') < 2^{j-1}\}$
  Initialize $\rho_i \leftarrow 0$ and increase ρ_i until $\sum_{i' \in W(i)} \rho_{i'} \geq 2^{j+1}$ or $\rho_i = \pi_i$
  if $\sum_{i' \in W(i)} \rho_{i'} \geq 2^{j+1}$ **then**
    Add i to Z and buy (i,z), i.e. $H \leftarrow H \cup \{(i,z)\}$
  else
    Pay penalty π_i
  end if
end while

Analysis. Let X be the set of terminals. For each buy terminal $z \in Z$, the algorithm incurs a cost of a_z. Thus, the total cost of the algorithm is $\sum_{z \in Z} a_z + \sum_{i \in P} \pi_i$. Our goal in the analysis is to show that this is at most $\mathrm{OPT}(T)$ for any HST embedding of (X,d).

First, we show that the cost of the algorithm is at most twice the total cost share $\sum_i \rho_i$.

Lemma 7.2. *We have that* $\sum_{z \in Z} a_z + \sum_{i \in P} \pi_i \leq 2 \sum_i \rho_i$.

Proof. The algorithm pays the penalty π_i for terminal i only if $\rho_i = \pi_i$ so $\sum_{i \in P} \pi_i \leq \sum_i \rho_i$. Let $Z_j \subseteq Z$ be the set of class-j buy terminals. Since $a_z \in [2^j, 2^{j+1})$ for $z \in Z_j$, we have $\sum_{z \in Z} a_z \leq \sum_j 2^{j+1} |Z_j|$. We now show that $2^{j+1}|Z_j| \leq \sum_{i \in X_j} \rho_i$ for each class j.

Fix a class j. The witness set $W(z)$ of $z \in Z_j$ satisfies the following properties: (1) $\sum_{i \in W(z)} \rho_i \geq 2^{j+1}$; (2) $W(z) \subseteq X_j$; (3) $d(i,z) < 2^{j-1}$ for $i \in W(z)$. The first implies that $2^{j+1}|Z_j| \leq \sum_{z \in Z_j} \sum_{i \in W(z)} \rho_i$. We claim that $d(z,z') \geq 2^j$ for $z, z' \in Z_j$. This completes the proof since together with the second and third properties, we have that the witness sets of Z_j are disjoint subsets of X_j and so $\sum_{z \in Z_j} \sum_{i \in W(z)} \rho_i \leq \sum_{i \in X_j} \rho_i$.

Now we prove the claim. Observe that H is the subgraph produced by the online greedy Steiner tree algorithm if it were run on the subsequence of buy terminals Z, and for each $z \in Z$ we have that a_z is exactly the distance from z to the nearest previously-arrived buy terminal. Thus, we can apply Lemma 3.2 and get that $d(z,z') \geq 2^j$ for any $z, z' \in Z_j$, proving the claim. Putting all of the above together, we get $\sum_{z \in Z} a_z \leq \sum_j \sum_{i \in X_j} \rho_i = \sum_i \rho_i$. This finishes the proof of the lemma. □

Next, we show that the total cost share is at most a constant times the cost of the optimal solution on any HST embedding.

Lemma 7.3. $\sum_i \rho_i \leq 8\,\mathrm{OPT}(T)$ *for any HST embedding* T *of* (X,d).

Proof. Let T be a HST embedding of (X,d). First, we lower bound $\mathrm{OPT}(T)$ in terms of T's cuts. Define $R_j \subseteq X_j$ to be the subset of class-j terminals i with $\rho_i > 0$. Consider a level-j cut $C_e \in \mathcal{C}_j(T)$ and associate the terminals $R_{j+1} \cap C_e$ to C_e. (Note that a terminal can only be associated to one cut.) By definition, $e \in E_j(T)$ and has length 2^{j-1}. If $r \notin C_e$ then e lies on the (i,r) path in T for each associated terminal $i \in R_{j+1} \cap C_e$ so the optimal solution on T either pays the penalty for each associated terminal at a cost of $\sum_{i \in R_{j+1} \cap C_e} \pi_i$ or buys e at a cost of 2^{j-1}. Since each terminal is associated to a unique cut, we have

$$\mathrm{OPT}(T) \geq \sum_j \sum_{C \in C_j(T): r \notin C} \min \left\{ \sum_{i \in R_{j+1} \cap C} \pi_i, 2^{j-1} \right\}.$$

We have $\sum_i \rho_i = \sum_j \sum_{i \in R_j} \rho_i$. For each terminal $i \in R_j$, we charge ρ_i to the level-$(j-1)$ cut $C \in \mathcal{C}_{j-1}$

containing i. Each level-j cut $C \in \mathcal{C}_j(T)$ is charged $\sum_{i \in R_{j+1} \cap C} \rho_i$. Thus, we have

$$\sum_i \rho_i = \sum_j \sum_{C \in \mathcal{C}_j(T)} \left(\sum_{i \in R_{j+1} \cap C} \rho_i \right).$$

Since $\rho_i \leq \pi_i$ for each terminal i, it suffices to prove the following claim: for each level-j cut $C \in \mathcal{C}_j(T)$, we have $\sum_{i \in R_{j+1} \cap C} \rho_i = 0$ if $r \in C$, and $\sum_{i \in R_{j+1} \cap C} \rho_i \leq 2^{j+2}$ if $r \notin C$. Suppose $r \in C$. Since $a_i < d(i,r) < 2^j$ for each $i \in C$, there cannot be any class-$(j+1)$ terminal in C and so $\sum_{i \in R_{j+1} \cap C} \rho_i = 0$. Now consider the case $r \notin C$. Suppose, towards a contradiction, that $\sum_{i \in R_{j+1} \cap C} \rho_i > 2^{j+2}$. Let i^* be the last-arriving terminal of $R_{j+1} \cap C$, i.e. i^* is the last-arriving class-$(j+1)$ terminal of C with $\rho_{i^*} > 0$. Since the diameter of C is less than 2^j, we have $W(i^*) \supseteq R_{j+1} \cap C$. We also have $\sum_{i \in R_{j+1} \cap C \setminus \{i^*\}} \rho_i < 2^{j+2}$ because otherwise the algorithm would not have increased ρ_{i^*}. But the algorithm increased ρ_{i^*} ensuring that $\sum_{i \in W(i^*)} \rho_i \leq 2^{j+1}$, contradicting the assumption that $\sum_{i \in R_{j+1} \cap C} \rho_i > 2^{j+1}$. This completes the proof of the claim and so we get $\sum_i \rho_i \leq 8\,\mathrm{OPT}(T)$. □

Now, Lemmas 7.2 and 7.3 imply that the cost of Algorithm 5 is at most $O(1)\,\mathrm{OPT}(T)$ for any HST embedding T of (X,d). Furthermore, by Theorem 2.1, there exists a HST embedding T^* such that $\mathrm{OPT}(T^*) \leq O(\log k)\,\mathrm{OPT}$. Thus, the algorithm is $O(\log k)$-competitive for prize-collecting Steiner tree and this proves Theorem 7.1.

Acknowledgements

The author is very grateful to Shuchi Chawla and Mohit Singh for helpful discussions and to the anonymous referees for their useful comments.

References

[1] Ajit Agrawal, Philip Klein, and R. Ravi. When trees collide: An approximation algorithm for the generalized steiner problem on networks. *SIAM J. Comput.*, 24(3):440–456, June 1995.

[2] Baruch Awerbuch and Yossi Azar. Buy-at-bulk network design. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 542–547, Oct 1997.

[3] Baruch Awerbuch, Yossi Azar, and Yair Bartal. On-line generalized Steiner problem. *Theoretical Computer Science*, 324(2-3), September 2004.

[4] Yair Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 184–193. IEEE Comput. Soc. Press, 1996.

[5] Yair Bartal. On approximating arbitrary metrices by tree metrics. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 161–168, New York, NY, USA, 1998. ACM.

[6] Piotr Berman and Chris Coulston. On-line algorithms for steiner tree problems (extended abstract). In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 344–353, New York, NY, USA, 1997. ACM.

[7] Friedrich Eisenbrand, Fabrizio Grandoni, Thomas Rothvoß, and Guido Schäfer. Approximating connected facility location problems via random facility sampling and core detouring. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1174–1183. Society for Industrial and Applied Mathematics, 2008.

[8] Jittat Fakcharoenphol, Satish Rao, and Kunal Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *Journal of Computer and System Sciences*, 69(3):485–497, November 2004.

[9] Dimitris Fotakis. A primal-dual algorithm for online non-uniform facility location. *Journal of Discrete Algorithms*, 5(1):141–148, 2007.

[10] Dimitris Fotakis. On the competitive ratio for online facility location. *Algorithmica*, 50(1):1–57, 2008.

[11] Michel X. Goemans and David P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.

[12] Anupam Gupta, Ravishankar Krishnaswamy, and R. Ravi. Online and stochastic survivable network design. *SIAM Journal on Computing*, 41(6):1649–1672, 2012.

[13] Anupam Gupta, Amit Kumar, Martin Pál, and Tim Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. *Journal of the ACM (JACM)*, 54(3):11, 2007.

[14] Makoto Imase and Bernard M. Waxman. Dynamic steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991.

[15] Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21(1):39–60, 2001.

[16] Adam Meyerson. Online facility location. In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 426–431. IEEE Comput. Soc, 2001.

[17] Jiawei Qian and David P. Williamson. An $O(\log n)$-competitive algorithm for online constrained forest problems. In Luca Aceto, Monika Henzinger, and Jiří Sgall, editors, *Automata, Languages and Programming*, volume 6755 of *Lecture Notes in Computer Science*, pages 37–48. Springer Berlin Heidelberg, 2011.

[18] Mário César San Felice, David P. Williamson, and Orlando Lee. The online connected facility location problem. In Alberto Pardo and Alfredo Viola, editors, *LATIN 2014: Theoretical Informatics*, volume 8392 of *Lecture Notes in Computer Science*, pages 574–585. Springer Berlin Heidelberg, 2014.

[19] Chaitanya Swamy and Amit Kumar. Primal–Dual Algorithms for Connected Facility Location Problems. *Algorithmica*, 40(4), September 2004.